

State of California - OCIO

Platform Domain Architecture

Version 0.01 (DRAFT)

September 9, 2009



[2009]

Table of Contents

Executive Overview.....	4
Technical Discussion of Current and Emerging Technologies.....	5
Platform Servers and Personal Computers.....	5
Software Partitioning and Virtual Infrastructure.....	5
Data Storage.....	7
Principles.....	13
Technical Standards.....	21
Introduction.....	21
Technology Components.....	23
Platforms.....	25
Computing Platforms Standards.....	25
Implementation Guidelines - Platforms.....	25
Best Practices for Platforms.....	25
Software.....	26
System Software Standards.....	26
Implementation Guidelines - System Software.....	26
Best Practices for Virtual Partitions.....	27
Data Storage.....	28
Data Storage Standards.....	28
Implementation Guidelines – Data Storage.....	28
Best Practices for Data Storage.....	30
Appendix One - Standard Desktop Configuration.....	32

History of Changes

Date	Modification
9/8/2009	Initial Draft by Lee Mosbrucker

DRAFT

Executive Overview

Mission Statement

The State of Connecticut Platform Architecture will identify technology hardware platforms and the related operating systems to support the State of Connecticut's current and future business requirements. As the State of Connecticut moves to a more centralized Information Technology enterprise, reductions in platforms and operating systems will reduce the total cost of ownership and accelerate the ability of IT to respond to the State of Connecticut's dynamic business needs.

Introduction and Background

Platform Architecture is a physical implementation within a logical application design. The platform decision for end users, applications, or databases is dependent on the overall application design in accordance with business requirements. Distributed computing client/server models focus on utilizing a variety of resources to make the end user more productive.

Adopting N-tier client/server application architecture requires that the software tiers be implemented on hardware tiers. Hardware tiers entail different kinds of computers performing different functions and exist to maximize the usefulness of various specialized hardware devices. Suitably designed applications incur little or no impact due to changes in platform strategy. Hardware dependencies should be minimal. In an N-tiered client/server environment, the mainframe becomes just another server.

Multiple software tiers may execute on a single or multiple hardware platforms. Pure mainframe applications utilize one platform for the user interface, program logic and data access. In a distributed computing environment, each facet may run on separate platforms. Moving to an N-tiered, distributed client/server application architecture introduces platform and operating system infrastructure and integration challenges. Interoperability across all platforms is crucial. Basing the State of Connecticut's platform architecture decisions on this document will promote the ability of software and hardware on different machines from different vendors to share data. It will also ensure easier integration by utilizing compatible technologies.

The Platform Architecture describes the platform requirements for supporting the hardware and operating systems that facilitate the implementation of a distributed computing environment. For both of the following sub-components, the principles in this domain help guide the evaluation, selection, design, construction, and implementation of the domain and its elements. The principles are focused upon the operating system and hardware selections with regards to the following:

Client

The client side of client/server is the interface to the application. Usually, a client runs on top of an operating system that provides a user interface to an application and access to distributed services. This includes handheld systems as well as notebook and desktop systems.

Server

The server provides services requested by clients. The server application runs on top of an operating system. Servers are often specialized by the type of function they perform on behalf of the client. They include application servers, database servers, collaborative technologies, middleware servers, server appliances and file and print servers.

Technical Discussion of Current and Emerging Technologies

Platform Servers and Personal Computers

Software Partitioning and Virtual Infrastructure

The concepts of partitioning and virtual infrastructure have been part of computer systems for over 15 years. These technologies were originally developed for mainframes to allow multiple operating systems or applications to share a single piece of hardware. IBM called these technologies LPARs (Logical Partitions) and VM (or Virtual Machine operating system). Since then the technologies have been implemented in mid-range hardware platforms (e.g., Sun Solaris, HP-UX/PA and IBM AIX/). More recently, the technology has been implemented in Intel Pentium and XEON based servers and workstations.

Partitioning

Partitioning can be accomplished through hardware and specialized operating system components (sometimes called logical partitioning or LPARs), or through software only (virtual partitions).

Partitioning has a number of advantages for reducing Total Cost of Ownership and increasing reliability and capability for disaster recovery:

- consolidates applications and infrastructure services running on diverse operating systems onto fewer, highly scalable servers
- facilitates lower level management of servers for more optimum performance levels
- helps to streamline testing and deployment
- promotes portability of consistent server images; and facilitates centralized management of servers.

Logical partitioning

Logical partitioning allows multiple applications or services to run simultaneously on a shared platform, with each application or service having its own instance of the underlying operating system. These instances can be of different revision levels. The applications or services also have their own instances of hardware resources (e.g., memory, disk or network ports). The partitions are isolated from one another, so that system anomalies or crashes cannot affect the other partitions. In addition, the resources of the system can be separately allocated a managed relative to each partition.

Virtual Partitioning

Virtual Partitioning uses a software layer or interface to accomplish what the hardware based logical partitioning does. This can be “above” the native operating system (e.g., VMware GSX or Microsoft Connectix), or can be it’s own operating system (e.g., VMware ESX or IBM VM). The software layer intercepts operating system calls for resources and translates them into a common set of interfaces. Thus each virtual partition sees its own view of the underlying hardware. (See Figure 1 at right).

Shared data storage (see next topic, can be used to transfer entire virtual partitions between physical partitions, or to be replaced by other virtual partitions with last current state intact. As in the case of logical partitions, resources can be allocated and managed between and among the virtual partitions.

Virtual Infrastructure

Virtual infrastructure (or virtualization) extends the concept of virtual partitions even further. Virtualization provides a “layer of abstraction” between the computing, storage and networking hardware, and the software that runs on it or utilizes it (see Figure 2). Virtual infrastructure treats hardware as a single pool of processing, storage and networking power to be allocated and de-allocated to various software services on the fly. Software management can then be separated from infrastructure management.

By creating a uniform virtual hardware platform, virtual infrastructure allows software and applications to be installed on or moved from one physical system to another without requiring reconfiguration of the operating system or applications. In advanced implementations, this movement can be accomplished without interruption of the application, and resources can be dynamically allocated or provisioned.

Storage Area Networks and Networked Attached Storage form the basis of extended virtual or logical partitions and infrastructure virtualization. These technologies are discussed in the next section.

Practical Uses for Partitioning

Beyond uses that reduce TCO, partitioning is a consideration when its two primary strengths are exploited, that is, sharing of resources and isolation of operating environments and applications from each other.

Example 1- Replacing older infrastructure

An agency has four older servers (for example Pentium 3, 1GHz) used for file and print serving, and several applications. These servers run a mix of Windows NT and Windows 2000 operating system and are generally running under 50% utilization. If partitioning were not used, updating these servers would mean buying four new servers (for example XEON 3.2 GHz). If any application needed NT to run, then NT would have to be ported to the new server, unless a rewrite of the software was undertaken by the agency. There would still be four separate boxes to manage, and each box would be substantially underutilized.

With the use of partitioning software (in this case VMware EX), one dual processor server could likely be used to replace all four boxes. This is because the processing power of a single XEON CPU could be allocated to replace two or even three Pentium 3 CPUs. The remaining CPU could be utilized to provide additional processing power, i.e., growth potential, for applications.

TCO would be enhanced and capital expenditures would be reduced. In addition there would be better utilization of local storage, or reduced connectivity expenditures if a SAN or NAS was used for storage (see next topic).

Example Two - development environment

Best practices for development environments call for separate servers for each to the logical and physical tiers of an N-tier environment. Without the use of partitioning, this would mean provisioning several physical servers, as well as a network infrastructure. With either virtual or hardware (LPAR) partitioning, the entire N-tier environment could be provisioned within a single server, and could utilize inter-partition communication facilities, thus avoiding the need for an external network infrastructure (unless one was desired).

Example Three - Rapid Deployment of Server “Images”

With the increasing use of low profile, rack mounted servers and blade servers, there is a growing need to be able to rapidly clone or replicate a given server image (operating system, drivers, applications, user profiles, etc.). This has application in several disaster recovery scenarios and in rapidly provisioning additional server capacity. Products such as VMware are particularly well suited to handling this task for Intel processor based systems.

Additional guidance on using VMware is provided in the Best Practices section of this document.

Data Storage

Data storage for servers and LANs can be categorized in one of three ways.

1. Directly Attached Storage (DASD), which is “dedicated” to an application or resource server. DASD is generally accessed through the file system of the server or workstation operating system (for example NTFS in Windows 2000 or XP). The OS file system is responsible for keeping track of the physical location of the data.
2. Networked Attached Storage (NAS), in which a “file server” uses a Local Area Network to connect application or resource servers to a large number of shared hard drives (DASD).
3. Storage Area Networks (SANs), which utilize specialized storage systems (called arrays) that have dedicated controllers and large shared pools of hard drives (DASD). Common practice is to utilize Fibre channel networks (or fabrics) to connect the storage systems to application or resource servers.

Generally there are six business objectives that can be achieved through the use of networked storage.

1. Consolidation of information resources and combining them for more effective management and security.
2. Collaboration or sharing of information thus enhancing reliability and accuracy, and reducing duplication.
3. Distribution of information for improved access, for back up and restoration, or for business continuity.

4. Minimizing the drain on processing capacity of the server when dealing with storage management, thus avoiding the degradation of its performance and investment (a feature of SANs).
5. De-coupling the replacement/upgrade lifecycles of servers from storage.
6. Providing incremental "as-needed" growth and configuration of very large pools of mass storage.

Network Attached Storage (NAS)

NAS storage systems generally use an IP based local area network to connect the storage resource to application servers or workstations. Files are accessed by name (from the application OS file system), these names are translated into physical locations by a dedicated "file server" in the NAS device. The file system of a NAS "file server" is responsible for keeping track of the physical location of the data rather than the OS of the application server or workstation. NAS devices are highly suitable for sharing of files.

Typically, NAS servers are used as a central repository for file and print servers (holding user folders and files, as well as spooling areas for printing), for control files (e.g., user profiles, directories, access lists, etc), and to store libraries of drive images, software updates and technical manuals. Besides centralizing access to data storage, NAS systems improve file access performance by offloading such processing from the application servers.

Storage Area Networks (SANs)

Unlike, NAS devices, SANs utilize block level access to data. This means the application OS knows where the data is located and "directly" accesses it. SANs utilize control processors, which can partition data and can utilize advanced RAID configurations to enhance both performance and reliability. The high-end SAN arrays (for example EMC CX 500/700 series), can simultaneously support multiple RAID protocols to fine tune the performance characteristics for different applications (e.g., RAID 0 or 1 for DBMS log files and RAID 5 for DBMS data tables). The controllers are also used for data duplication or for backup and restore procedures, off-loading such duties from the processors in the server. The controllers are also used to facilitate SAN to SAN connectivity for disaster recovery or data distribution.

Generally, SANs utilize fibre channel technology (or fabrics) to connect servers or workstations to the SAN device and to connect SAN controllers to the SCSI disk drives (at least in high performance SANs). SANs can also be implemented with ATA type disk drives, which offer higher storage densities at the cost of reduced performance. Sometimes low-end SANs might be implemented with IP networks (though at lower throughput).

Typically SANs are employed where consolidation of storage is of high importance, as is scalability and throughput. SANs also provide higher levels of availability for mission critical applications and large scale DBMS systems involving all forms of storage for data, files, images, videos, and other electronic media.

Fibre Channel Networks (Fabrics)

Figure 3 Fibre Channel Network Topologies.

While not a covered technology within the Platform Domain Architecture, Fibre Channel networks (generally called fabrics), have an important part to play in the design and implementation of a Storage Area Network. Fibre networks utilize a high speed serial protocol and dedicated channel to transmit frames of data between devices. This contrasts with Ethernet (or IP networks) which transmits packets over a shared channel, which can lead to contention or collisions and thus reduced throughput with heavy loads. In addition, all network processing is confined to the Host Bus Adapter (HBA) thus reducing CPU overhead to less than 5% (compared to some Ethernet implementation which can exceed 30% overhead).

While Fibre networks can be point to point or loop topologies, the common practice is to utilize high speed switches to create a "fabric" (see Figure 3 at right). The characteristics of the Fibre channel switch are such that there is equal performance between any pair of ports.

As Fibre networks generally run at 1 Gbit/sec or 2 Gbit/sec, this means data can move between a server and array or array and backup at up to 200 MByte/sec. This is fast enough to allow for dual paths to be created between devices to allow for fail-over should one path become non-operational.

Smaller Fibre channel networks might use a single switch to handle connections. Larger networks utilize what are called edge switches which are interconnected by core or director switches in what can best be described as a mesh. This is shown below in Figure 4. This type of network topology supports dual paths or the fibre channel connections (the blue and black lines). The general approach of core and edge switches (or storage devices) is the preferred approach for SAN designs as it facilitates flexibility in designing pathways for performance and in scalability.

In addition to Fibre Channel network fabric, high-end SANs also utilize Fibre channel connections between the individual drives and the SAN controller. This enhances performance and allows the utilization of "zones" or dedicated pathways between parts of the array to external devices that are connected through the Fiber fabric.

General Implementation Best Practices

Multiple Tiers of Storage

The basic design of a SAN takes into account three tiers of storage. These classes allow for the matching of data to a storage technology that has an appropriate level of performance and availability. This gives a lower cost of ownership without any negative impact on business. These three classes or tiers of data storage are:

- A first tier on fibre channel storage arrays — data with the highest access requirements;
- A second tier on SATA disks within a storage array (or possibly within a NAS) — near-online or infrequently accessed data; this serves as the short term business backup and as the staging area for the third tier. This tier makes use of special software executing in the control unit in the SAN array.
- A third tier of storage on tape or even on optical media — archived data and long term backup copies; these copies are stored off-site in a secure, resilient location. (A

new alternative is to use specialized, object based storage arrays for archiving). DOIT has historically implemented tier one and tier three (that is direct backup from disk to tape). With the deployment of new SAN arrays, DOIT will begin to utilize a three tier model.

Multi-Pathing

As discussed previously, Fibre Channel Fabrics often use multi-pathing for interconnections. This means two or more Host Bus Adapters (HBAs) in each device attached to the fabric, terminating fibre channels in different switches and software installed in each device to allow for fail-over to an alternate path (or paths) in the event of the failure of a fabric component. DOIT routinely installs dual HBAs in servers and runs multiple paths to storage arrays and to the components of the tape backup system.

Storage Replication

Storage replication copies data from one SAN to a secondary SAN array (which could be located at a second data center) on either an asynchronous or synchronous (mirrored) basis. This is accomplished by software executing in the control unit of the SAN array. Either fibre channel or

IP networks can be used as the transport mechanism. The primary use of replication within the State will be for business continuity and disaster recovery.

Planning

HP recommends¹, at a minimum, that you document the following before beginning the actual implementation:

1. Topology Map—Shows the logical SAN topology and fabric interconnect scheme; conveys the overall design from a strategic standpoint, and can also serve to convey how future growth and technological advances will be accommodated.
2. Configuration Layout—Shows the physical layout of the entire implementation. More detailed than the topology map, the layout is used during implementation to verify the correct connectivity. This is also extremely helpful if troubleshooting is required in later phases.
3. Storage Map—Defines the storage system arrangement and configuration in the SAN, and storage settings such as SSP and RAID levels. This map effectively defines how all of the storage is configured in the SAN.
4. Zoning Map—Defines the inter-node communication access within the SAN. This map defines which nodes or user ports are allowed to communicate with each other in the SAN.

In addition, you should consider each of these items during the planning phase:

- Deployment Strategy: You can choose to deploy separate smaller SANs or SAN Islands with the idea of increasing capacity by growing the SANs independently or by interconnecting the independent SANs in the future. Smaller SANs are easier to construct, larger SANs offer economies of scale from an operational standpoint, but take longer and are more complex to build.
- Topology Design: Consider the topology design compared to the ease of migrating to another, higher capacity design. In most cases this can be accommodated; however, it is always preferable to choose an initial design that can grow, without the need to transition to a different topology.
- Experience Level: If you are just beginning deployment of SAN technology, consider starting with a smaller implementation. As you gain experience, deploy larger SANs.
- SAN Management Strategy: Define the management strategy and the specific tools that you will utilize to manage your SAN.

1 SAN Design Reference Guide AA-RU5ZD-TE (November 2003), pp 236-237.

Backup/Restore

The discussion of multiple tiers of storage introduced the issues and business needs of backup for disaster recovery and business continuity and disaster recovery and archiving for meeting legislative or regulatory requirements for data retention, as well as moving inactive data off of production storage.

A backup is a copy of files, directory, data records, logs, etc. to a second medium (a disk or tape) as a precaution in case the first medium fails or becomes corrupted. As such, backup is an important component of disaster recovery. The backup copy can take any number of forms including:

- An exact copy of a physical disk structure. Programs such as Ghost are used for this purpose. An image must be restored to an identical hardware configuration. This approach is often used to backup the system files of computer system. (Note: this is called an "image copy")
- A file by file copy of the logical directory structure on a disk. If another disk is the target then the copy process replicates the directory and file structure on the target device. This type of backup can generally be read by any compatible operating system. (Note: this is sometimes called a "mirror backup")
- A copy of just the files on a disk along with metadata information on where the file belongs. This copy can be to another disk or to magnetic tape and it might be compressed to save space. Generally a specialized program is used to create the backup and to read or restore the data.
- A compressed file backup which creates a "large" file containing the file contents and associated metadata. This requires a specialized program or utility. WinZip™, popular program for personal computers, is an good example of such a program. Microsoft Windows backup utility is another. While disk drives and tape are typical target devices, optical storage (CD-ROM or DVD) are becoming common as a target device.

In addition to the form of the actual backup, there are several approaches that can be employed when making backup copies:

- Full backup — which is literally making a complete copy of the desired material, regardless of whether or not it has changed. This type of back-up is very time consuming and consumes the most space, however, restores can be very fast.

- Differential backup — the backup of only the data objects or files that have been modified since the last cumulative incremental backup or full backup. Differential incremental backups run fast and take the least amount of space.

- Cumulative incremental backup — the backup of the data objects or files that have been modified since the last full backup, this encompasses all prior incremental backups since the last full backup. In extreme cases, this becomes "incremental forever", in which only an initial full backup is made, everything else is differential.

Incremental backup programs (for example IBM Tivoli Storage Manager or TSM) utilize specialized catalog or log files to track which files have been changed rather than relying on the source system operating system to track changes.

The advantage of lower backup times for any incremental backup approach comes with a price: increased restore time. When restoring from incremental backup, one needs the most recent full backup as well as every incremental backup made since the last full backup. This can take a lot of time, and can be very complex if the backup has been spread (stripped) across multiple tapes to increase throughput.

In addition, the size of the accumulated differential backup files can exceed the size of a full backup.

Archiving

Archiving is focused on long term retention of data and information on third-tier, offline storage media. This media can be magnetic tapes, optical disk (Read/Write or Write Once) and COLD2, microfilm, or specialized storage arrays (for example Content Addressable Storage3 as in the EMC Centera product line). If required by regulations or statutes, only optical storage or microfilm can provide truly long term, immutable, authenticated storage of archived data.

Approaches to archiving vary with the type of information that needs to be stored (e.g., electronic documents, digital X-rays, check images, movies, e-mail, and broadcast content). For instance MS Word documents can be archived in original digital file form (in compressed form), converted to a PDF document for subsequent storage or processed by a COLD system. Printed output can be stored in a COLD system or on microfilm. Microsoft Exchange presents its' own set of archiving problems, as a distinction can be made between retrieving mail at the mail store, mailbox or individual e-mail or calendar entry level. While tape or optical disk can be the storage medium for the archive, Content Addressable Storage systems can also be used, if only for 3 to 5 year storage periods.

2 Computer Output to Laser Disk - (COLD) The capture of large reports or documents on optical media such that sections are accessible as individual documents. Most cold systems utilize compression to save space.
COLD

systems maintain indices on stored contents to allow for rapid retrieval. A successor technology to "COM" (Computer Output on Microfilm).

3 Content addressable storage is a storage system that uses the content of the data (or part of the contents) to identify the locator for the information. Typically, this is implemented by using a polynomial algorithm on the data to resolve to a unique ID. That ID, sometimes called a signature, is then associated with metadata as to the actual location of the data. Only the ID is then available to the application, or user if you will, to be able to access that data. One unique aspect is that if the exact same data is written again, it will resolve to the save identity so that no duplicate data will actually be stored.

Principles

The principles listed below provide guidelines for the design and selection of platform technology components that will support n-tier solutions deployed across the enterprise. These principles are general and apply to both server and client-side solutions.

Principle 1: Architecture Management

The State of Connecticut's enterprise-wide technical architecture must be unified and have a planned evolution that is governed across the enterprise.

Rationale

Without a unified approach, there will be multiple, and possibly conflicting, architectures.

- Good change requires collaboration and collective planning across the enterprise.
- Architecture must be well thought out.
- Governance will be simplified.

Implications

Normal evolution will require prioritization and re-prioritization across all platform and operating systems initiatives.

- Dependencies must be maintained.
- The architecture must be continually re-examined and refreshed.
- Short-term results vs. long term impact must be constantly considered.
- Establishing enterprise architecture takes time and involves a lot of change.
- Architecture decisions must be verified with the Enterprise IT Management for feasibility, migration paths and implementation schedules.

Principle 2: Architecture Compliance

Architecture support and review structures shall be used to ensure that the integrity of the architecture is maintained as systems and infrastructure are acquired, developed and enhanced.

Rationale

To realize the benefits of a standards-based enterprise architecture, all information technology investments must ensure compliance with the established IT architecture.

For maximum impact, review should begin as early in the solution planning process as possible “If you are going to talk the talk, then you must be willing to walk the walk.”

Implications

A structured project level review process will be needed to ensure that information systems comply with the IT Architecture and related standards.

Processes incorporating the principles of this (technical) architecture must be developed for all application procurement, development, design, and management activities.

This compliance process must allow for the introduction of new technology and standards.

Conceptual Architecture and Technical Domain principles should be used as evaluation criteria for purchasing as well as developing software.

Principle 3: Ensure Security, Confidentiality and Privacy

IT systems should be implemented in adherence with all security, confidentiality and privacy policies and applicable statutes.

Rationale

- Helps to safeguard confidential and proprietary information
- Enhances public trust
- Enhances the proper stewardship over public information
- Helps to ensure the integrity of the information

Implications

- Need to identify, publish and keep the applicable policies current.
- Need to monitor compliance to policies.
- Must make the requirements for security, confidentiality and privacy clear to everyone.
- Education on issues of privacy and confidentiality must become a routine part of

normal business processes.

- Host level security must be enforced by the responsible System Administrator to
- prevent unauthorized access the system.

Principle 4: Reduce Integration Complexity

The enterprise architecture must reduce integration complexity to the greatest extent possible.

Rationale

- Increases the ability of the enterprise to adapt and change.
- Reduces product and support costs

Implications

- Decreases the number of vendors, products, and configurations in the State's environment.
- Must maintain configuration discipline.
- Will sacrifice performance and functionality in some instances.
- Will rely on components supplied by vendors.
- Determination of "the greatest extent possible" includes consideration of how reducing complexity can negatively impact providing critical client services.
- New product selections must weigh current and established platform standards when choosing new product or point solutions.

Principle 5: Integration

Systems must be designed, acquired, developed, or enhanced such that data and processes can be shared and integrated across the enterprise and with Connecticut's business partners.

Rationale

Increase efficiency while better serving our customers (e.g., the public, agencies, etc.).

Redundant systems cause higher support costs.

Ensures more accurate information, with a more familiar look and feel.

Integration leads to better decision making and accountability.

Implication

IT staff will need to consider the impacts on an enterprise wide scale when designing applications.

We will need new tools and training for their proper use.

Will need a method for identifying data and processes that need integration, when integration should take place, whom should have access to the data, and cost Rationale for integration.

Will need a “coordinator” that can maintain and arbitrate a common set of domain tables, data definitions, and processes across the organization.

Over integration can lead to difficult data management and inefficient processes.

Principle 6: Total Cost of Ownership

Adopt a total cost of ownership model for applications and technologies which balances the costs of development, support, training, disaster recovery and retirement against the costs of flexibility, scalability, ease of use, and reduction of integration complexity.

Rationale

- Consideration of all costs associated with a system over its entire life span will result in significantly more cost effective system choices.
- Enables improved planning and budget decision-making.
- Reduces the IT skills required for support of obsolete systems or old standards.
- Simplifies the IT environment.
- Leads to higher quality solutions.

Implications

- The State budget process needs to accommodate Total Cost of Ownership of a system over a longer timeframe than current budgeting models.
- Will require looking closely at technical and user training costs especially when making platform or major software upgrades during the lifetime of the system.
- Requires designers and developers to take a systemic view.
- Need to selectively sub-optimize individual IT components.
- Need to develop a cost of ownership model.
- Need to ensure coordinated retirements of systems.

Principle 7: Minimize Platform Configurations

Create a small number of consistent configurations for deployment across the enterprise.

Rationale

- Reducing uniqueness in product selection and standardization reduces support and maintenance costs, and simplifies training and skills transfer.
- The cost of IT personnel is increasing and the cost of hardware is decreasing rapidly.
- This is the most efficient approach to enterprise-wide infrastructure configuration and maintenance.

- By constantly ‘tweaking’ the performance of an individual server or desktop computer, a multitude of unique configurations is created, thus increasing support and maintenance costs.

Implications

- Increased initial capital investment
- Deploy applications on uniformly configured servers (“If in doubt, use the bigger Box”)
- Plan to replace multiple, non-standard, configurations with a small number of consistent configurations.
- Plan for the regular replacement of platform components to ensure the retirement of obsolete and unique configurations
- Limits product choice and vendor selection when developing new applications

Principle 8: Basic Information Services

A standardized set of basic information services (e.g., email, voicemail, e-forms, user training) will be provided to all employees.

Rationale

- Increases productivity
- Reduces costs of maintenance
- Provides the basis for multi-agency or statewide business initiatives
- Provides for universal employee access to information

Implications

- Basic services definition needs to be created and regularly reviewed
- May increase “one-time” costs to upgrade to the minimum service level
- Training must be provided to all users of basic services
- Increase the need for support staff as more advanced technology is rolled out. (Helpdesk, System Administrator, networking and security)

Principle 9: Shared Components Using an N-Tier Model

Applications, systems and infrastructure will employ reusable components across the enterprise, using an n-tier model.

Rationale

- You can make significant changes to a component of a system, such as changing from a Windows client to a web-browser client, without changing the rest of the system
- Enables simplification of the environment and geographical independence of servers
- Takes advantage of modular off-the-shelf components
- Reuse will lower costs and maintenance efforts
- Allows for leveraging skills across the enterprise
- Must take advantage of the least technically complex solutions that meets the business needs of the project while maintaining flexibility within other n-tier principles to reduce system management complexity

Implications

Component management must become a core competency.

Requires the development a culture of reuse.

Reusable components must be platform independent.

Physical configuration standards must be established.

Design reviews become crucial.

Application systems must be highly modularized without making components too small or too simple to do useful “work”.

Principle 10: Mainstream Technologies

IT solutions will use industry-proven, mainstream technologies.

Rationale

Avoids dependence on weak vendors.

Reduces risk.

Ensures robust product support.

Enables greater use of commercial-off-the-shelf solutions.

Implications

Need to establish criteria for vendor selection and performance measurement.

Need to establish criteria to identify the weak vendors and poor technology solutions.

Requires migration away from existing weak products in the technology portfolio

Principle 11: Industry Standards

Priority will be given to products adhering to industry standards and open architecture.

Currently open and industry standards conflict and care must be taken to remain flexible and

avoid being locked into proprietary solution during this period of rapid change in the IT field.

Rationale

Avoids dependence on weak vendors.

Reduces risks.

Ensures robust product support.

DRAFT

Enables greater use of Commercial-off-the-Shelf solutions.

Allows flexibility and adaptability in product replacement.

Implications

Requires a culture shift.

Need to establish criteria to identify standards and the products using them.

IT organizations will need to determine how they will transition to this mode.

Migration away from systems and products with fading market share and viability.

Principle 12: Disaster Recovery / Business Continuity

An assessment of business recovery requirements is mandatory when acquiring, developing, enhancing or outsourcing systems. Based on that assessment, appropriate disaster recovery and business continuity planning, design and testing will take place.

Rationale

Due to factors such as the Internet and Y2K, customers and partners have heightened awareness of systems availability.

The pressure to maintain availability will increase in importance. Any significant visible loss of system stability could negatively impact our image.

Continuation of business activities without IT is becoming harder.

Application systems and data are valuable State assets that must be protected.

Implications

Systems will need to be categorized according to business recovery needs (e.g. business critical, non-critical, not required).

Alternate computing capabilities need to be in place.

Systems should be designed with fault tolerance and recovery in mind.

Plans for work site recovery will need to be in place.

Costs may be higher.

Principle 13: Scalability

The underlying technology infrastructure and applications must be scalable in size, capacity, and functionality to meet changing business and technical requirements.

Rationale

Reduces Total Cost of Ownership by reducing the amount of application and platform changes needed to respond to increasing or decreasing demand on the system.

Encourages reuse.

Leverages the continuing decline in hardware costs.

Implications

Scalability must be reviewed for both “upward” and “downward” capability.

May increase initial costs of development and deployment.

Will reduce some solution choices.

Technical Standards

Introduction

The Platform Architecture Principles illustrate the need to reduce the hardware and software configurations of platforms and operating systems supported by the State of CT’s Information Technology Department to increase flexibility, reduce cost, leverage, increase reuse and build on existing skills sets. The standards reflect the current state of the infrastructure at the State of Connecticut. These standards are based on a point in time snapshot of the existing State of CT IT infrastructure. These standards will provide a level-set foundation for future technology investments that assist in meeting the principles. These standards will be used as the State of CT’s Information Technology Systems mature and should be viewed as a migration path for existing platforms and as the target platforms for new systems. Platform technology will be periodically evaluated for a possible refresh of the existing standards.

NOTE: This document is not intended to rule out point solutions or niche players in certain technical areas, but to provide greater consistency across the enterprise. Research/experimental operating systems and platforms are not regulated by this document.

We are using the following categories to classify the technology standards and components.

Obsolete Standards

It is highly likely that these standards or products, while still in use, will not be supported by the vendor (industry, manufacturer, etc.) in the future. Some products and standards have already reached the non-supported state. Plans should be developed by the agencies or the State to rapidly phase out and replace them with strategic standards or products. No development should be undertaken using these standards or products by either the agencies or the State.

Transitional Standards

These are standards or products in which an agency or the State has a substantial investment or deployment. These standards and products are currently supported by DOIT, the agencies, or the vendor (industry, manufacturer, etc.). However, agencies should undertake development using these standards or products only if there are no suitable alternatives that are categorized as strategic. Plans should be developed by the agencies or the State to move from transitional to strategic standards or products as soon as practical. In addition, the State should not use these standards or products for development.

Note: many older versions of strategic standards or products fall into this category, even if not specifically listed in a domain architecture document.

Strategic Standards

These are the standards and products selected by the state for development or acquisition, and for replacement of obsolete or transitional standards or products. (Strategic means a three to four year planning horizon.) When more than one similar strategic standard or product is specified for a technology category, there may be a preference for use in statewide or multi-agency development. These preferred standards and products are indicated where appropriate.

Note: some strategic products may be in “pilot testing” evaluation to determine implementation issues and guidelines. Pilot testing must be successfully completed prior to full deployment by the agencies or the State.

Research / Emerging Standards

This category represents proposed strategic standards and products that are in advanced stages of development and that should be evaluated by the State. Some of these standards or products may already be undergoing “hands-on” evaluation. Others will need to be tracked and evaluated over the next 6 to 18 months.

Technology Components

General Technical and Product Standards

Specific version numbers will be found in Table 1 Computing Platforms or in Table 2 System Software.

- Standards 1:** The server operating systems standards for enterprise and application servers will be IBM Z/OS, Sun Solaris, HP UX and Microsoft Windows Server 2000/2003.
- Standards 2:** The hardware platforms for UNIX (e.g. SUN or HP/UX) and Windows 2000/2003 servers are either RISC based or Intel XEON based.
- Standards 3:** The server operating system standards to support NOS (network operating systems) is Windows 2000 and Windows 2003.
- Standards 4:** The operating system standards for "Dedicated Infrastructure Appliances" are either Linux or FreeBSD Unix. (Note: version is unique to manufacturer of the appliance).
- Standards 5:** The operating standard for "Dedicated Infrastructure Servers" will be either one of the State standard operating systems (Solaris, HP-UX, Windows Server 2003) or RedHat Enterprise Linux ES.
- Standards 6:** The hardware platforms for Windows 2000/2003 NOS are either DELL PowerEdge series or HP/Compaq ProLiant DL series. >
- Standards 7:** The standard server platform for the support of "thin clients" will be a combination of Citrix Terminal Services and Microsoft terminal server.
- Standards 8:** Workstations should be configured to reflect the intended use (e.g., CAD, GIS, development, testing). The platform and operating system are to be selected from products similar to server platforms and operating systems.
- Standards 9:** The standard desktop personal computer for State agencies is based on the Intel Pentium 4 (with hyper-threading) CPU and Windows XP operating system. See Product Standards table and implementation guidelines below.
- Standards 10:** The standard Storage Area Network (SAN) for State agencies is the EMC Clariion product line.. See Product Standards table and implementation guidelines below.

Rationale:

These operating systems are considered to be "open" or have the ability leverage open systems functionality. These operating systems allow flexibility and scalability across the enterprise. Consistent configuration eases the burden of maintenance, may reduce the purchase costs, and will greatly simplify migration to emerging standards in the future. Application size, complexity, migration path, system management and total cost of ownership will determine which platform will be deemed preferred.

Comments on changes to the Tables of Product Standards

The original product standards tables combined hardware and operating systems as the platform entry in the tables. This has been changed in the current version to separate hardware from software. The rationale for this is straightforward: platform technology in certain categories will change more quickly than the operating system. In other cases, while the specific model of platform is still "strategic" the operating system has moved on to a newer version with better security and operational features. Separating hardware from software allows for greater flexibility.

Another important change is the elimination of products that were listed as "obsolete" in Version 1 of this document that was published in January 2001. Another important change is the elimination of certain LAN Server and SUN products that are no longer manufactured, have no real manufacturer support or are no longer used within State agencies.

DRAFT

Platforms

Computing Platforms Standards

TABLE HERE

Implementation Guidelines - Platforms

LAN Server Hardware Platform

DELL PowerEdge 2650 Series (Intel XEON) servers are the preferred platform for LAN and entry level application servers. This is the de facto standard at the DOIT data center. Agencies not currently utilizing HP Proliant DL series servers will acquire DELL PowerEdge 2650 servers.

Agencies currently who have deployed HP Proliant DL series servers may continue to do so.

Personal Computers

The PC subcommittee of the Platform Domain Team has established a standard configuration for desktop (personal) computers. . The standard configuration uses embedded networking and graphics capability to minimize support requirements.

The current list of acceptable brands is: -DELL, HP, IBM, and Gateway.

Appendix One - Standard Desktop Configuration (Page 29) lists the components and specifications of the current standard configuration.

Best Practices for Platforms

Best Practice 1: Major applications should be placed on uniformly configured servers to make overall maintenance, support and recovery less expensive. New major applications should be written for the EWTA standard platforms.

Best Practice 2: When considering new servers or upgrades take into consideration future growth when sizing the platforms.

Best Practice 3: An effort should be made to come up with the uniform workstation configurations to reduce complexity and administration costs.

Best Practice 4: Build in the ability to securely manage remote servers and clients to reduce system administration costs.

Best Practice 5: Weigh business needs vs. system administration costs when upgrading or implementing new systems and establish a balance between business goals/flexibility and system management ease.

Best Practice 6: Both Win/Tel server and workstation hardware should be chosen from the top tier vendors. An effort should be made to reduce choices to agree with the EWTA platform principles.

Software

System Software Standards

TABLE HERE

Implementation Guidelines - System Software

Desktop Operating Systems

Consistent with State IT Policy, agencies are required to acquire Windows XP licenses with any new PC purchase or when upgrading from Windows NT, Windows 98 or Windows ME. They may install (through a retro license right) Windows 2000 for compatibility reasons, but they should preferentially install Windows XP.

LAN Operating Systems

Agencies may install either Windows 2000 or Windows 2003 as a LAN server operating system.

Dedicated Infrastructure Servers

RedHat Enterprise Linux ES in a production environment will only be used for DNS servers, FTP servers or Proxy servers. Other State standard operating systems may also be used with the decision being based on initial cost, support TCO, scalability and the appropriateness of the choice to the business requirements. The operating system for other types of dedicated infrastructure servers will be determined on a case by case basis.

RedHat Enterprise Linux ES will be acquired with a minimum of 4 hour response time 12x5 support from the publisher (RedHat).

RedHat Enterprise Linux ES is not a State standard for more generalized infrastructure servers in a production environment (for example web servers or DBMS servers), nor for general applications in a production environment.

Best Practices for Virtual Partitions

Using VMware:

Typical Situation for using VMware

Multiple file-and-print or other irregularly used servers; typical loads of 25-35% CPU utilization;

Older servers utilizing Pentium 2 or Pentium 3 CPU technology or early (1GHz) edition XEON CPUs; generally this class of server is 3 to 4 years old and may be targeted for refresh or replacement.

Replacement with a new, single server should occur at the planned refresh or replacement.

N-Tier application architectures that require physical partitioning, but that are not performance constrained.

Cost savings on implementing development and test environments (not only through virtual partitions, but also by "exchanging" entire operating environments within a partition).

Isolated DNS servers (DNS servers are used to resolve an Internet address) for highly secure applications such as those used by Public Safety or CJIS agencies. (Note: DNS servers represent a very light server load, but require separate Network Interfaces; VMware can effectively manage isolating communication channels and other physical resources).

Applications that need to be isolated from one another as their O/S requirements or versions conflict with one another.

Situations to be avoided

Applications that require more than one CPU for adequate performance.

Relatively new servers utilizing faster XEON processors (1.5 to 2 GHz).

Migration Considerations

P2V should be used for transferring server environments and applications to an ESX virtual partition, when originating environments are Windows NT 4, SP6 or Windows 2000.

Older Windows NT environments need to be evaluated to understand the costs of updating the environment (applications sometimes "break" when this is done) vs. costs of recreating the environment and reinstalling or migrating the applications. Expertise in the operating system environment is a prime requirement for moving

pre NT 4, SP6 applications.

- Typical Server Platform(s)

Dual or quad CPU (XEON 3.0 or 3.06GHz) rack mounted servers (ideal) or recent, but slower XEON based servers (depends on performance requirements).

RAID configured multiple drives, NAS or, when hosted at DOIT, connection to the SAN.

Minimum 4 GB of RAM

Data Storage

Data Storage Standards

TABLE HERE

Implementation Guidelines – Data Storage

Attached Storage

Storage Area Network technologies allow a “virtualization” of storage throughout the State on behalf of all agencies. The most significant cost effectiveness is achieved when the storage needs are aggregated into a single statewide SAN to facilitate storage on demand, regardless of where, when, and how much is needed. Products chosen should be configured to allow for both central and distributed administration and management.

Factors that need to be addressed in using SAN technologies

Network bandwidth to the server and the agency site.

Agency needs for backups, copies, disaster recovery, and replication.

Physically installing the storage arrays in data center class climates.

Agency ability to install and configure the Storage Area Network or Storage Array.

Agency collaboration with DOIT on both central and distributed administration and management.

Factors that need to be addressed in considering Direct Attach Storage (DAS) vs. Networked Attached Storage (NAS) technologies

Configurations of Direct Attached Storage (DAS) and Networked Attached Storage (NAS) combined with Storage Area Network arrays are often the most complex and require careful consideration. Such assessments often include:

Existing DAS and possible upgrades to that installed base;

Bandwidth to the agency site because a NAS unit can be anywhere;

Need for services that offer sophisticated backup, replication, and snapshot capabilities;

Use of NAS as a file interface to the SAN array with its block based efficiency;

Need to share a single file real time among participating servers (cluster) and the need to store many files or make a single file copy available to many clients (file server);

Need for remote administration and backup/recovery; and

Combining NAS with SAN and balancing the costs, administration, performance, bandwidth and mean time to failure.

Best Practices for Data Storage

Best Practice 7: Incorporate a “data life cycle management” practice into State and agency systems development project methodology to assist the agencies with defining their needs for data copies, retention, redundancy, privacy, HIPPA, FOI and other unique factored requirements

Best Practice 8: Establish pragmatic approaches for file and data conversions to this technology from servers to expand the life of servers and avoid costly storage upgrades

Best Practice 9: Design and implement a portfolio of product and service offerings to meet agency storage needs with a variety of technologies⁴. This portfolio will enable the choice of the most cost effective solution for any particular need. Stored videos do not use the same technology as a production transactional data file, for example.

Best Practice 10: Define a balanced server and storage hardware investment strategy to achieve an optimal outcome financially and operationally. This also must take into consideration the network sizing and skilled staff deployment.

Best Practice 11: Until such time as a singular integrated solution is available; utilize the Storage Area Network and Storage Array management software provided by the manufacturer of the SAN or Storage Array.

Best Practice 12: Design large SANs to have two independent fabrics. Each fabric operates independently, and the failure of one fabric does not cause a complete loss of SAN communication.

The two fabrics should be similar in size and topology. This minimizes the risk of asymmetrical performance under certain workloads, and minimizes the total cost of the SAN.

⁴ While this best practice has a DOIT focus, it would apply within an agency as well with respect to having more than one solution to agency data storage needs.

Best Practice 13: Use a “core plus edge switch” combined with the dual independent fabrics approach to facilitate growth and configuration changes for SAN fabric design. The initial design should include spare ports on the core to support the future addition of edge switches.

Best Practice 14: When setting up zoning, use meaningful names for zones and zone aliases and be consistent with the naming convention throughout the fabric.

Best Practice 15: During the configuration phase, it is important to document the details of the actual physical configuration. This documentation should include items such as:

- Recording the WWN of all nodes and devices and identify where they physically reside.
- Define a system for cable labeling..
- Labeling switches using a relevant naming scheme particular to the topology.

- Defining Zones using a zoning map (or equivalent) to configure zones.

DRAFT

Appendix One - Standard Desktop Configuration

TABLE HERE

DRAFT